A Novel No-reference Video Quality Metric for Evaluating Temporal Jerkiness due to Frame Freezing

Yuanyi Xue, *Student Member, IEEE,* Beril Erkin, and Yao Wang, *Fellow, IEEE*

arXiv:1411.1705v1 [cs.MM] 5 Nov 2014

***Abstract*—In this work, we propose a novel no-reference (NR) video quality metric that evaluates the impact of frame freezing due to either packet loss or late arrival. Our metric uses a trained neural network acting on features that are chosen to capture the impact of frame freezing on the perceived quality. The considered features include the number of freezes, freeze duration statistics, inter-freeze distance statistics, frame difference before and after the freeze, normal frame difference, and the ratio of them. We use the neural network to find the mapping between features and subjective test scores. We optimize the network structure and the feature selection through a cross validation procedure, using training samples extracted from both VQEG and LIVE video databases. The resulting feature set and network structure yields accurate quality prediction for both the training data containing 54 test videos and a separate testing dataset including 14 videos, with Pearson Correlation Coefficients greater than 0.9 and 0.8 for the training set and the testing set, respectively. Our proposed metric has low complexity and could be utilized in a system with realtime processing constraint.**

***Index Terms*—video quality metric, neural network, packet loss, temporal jerkiness.**

I. Introduction

FRAME freezing and the consequent temporal jerkiness is a commonly observed artifact in Internet video applications due to both packet losses and delays. Depending on the delay allowance of the underlying video applications, there are two types of frame freezing artifacts. In applications with stringent low delay requirement, (such as video conferencing or live streaming), any frame that is not completely received by its display deadline is considered lost, and the receiver chooses a certain error concealment method to recover the frame. A common and popular error concealment method is to use its previous frame that was correctly received. The subsequent frame, even if correctly received, will have decoding error, if it is predicatively coded using the previous frame. To avoid such error propagation problem, all subsequent frames after a lost frame are also replaced by the same last correctly received frame, until the next intra-frame is received. We call the resulting artifact "frame freeze due to packet loss". In applications allowing more elastic delay, such as streaming of pre-coded video, when a frame does not arrive past its display deadline, the receiver continuously displays the previous frame, until the actual new frame arrives. We call this artifact "frame freeze due to packet delay". Both artifacts manifest as temporal jerkiness on the received video. In this work, we propose a metric that quantitatively measures this temporal jerkiness and its relation to the subjective quality.

The proposed metric is a no-reference (NR) metric, in that it evaluates the jerkiness based on the degraded video only, rather than by comparing the degraded video with the original pristine video. No-reference metrics are important for quality assessment in real applications, as the pristine video is often not available at the receiver. There are several related works for detecting and measuring the frame freeze. In [1], the authors proposed an NR metric for measuring the quality impact of frame freeze, based on the profiles of the duration of each freeze and the number of freezes. However, the features they considered are not dependent on the video content, which is undesirable, since for different video characteristics, the same freeze pattern could have different impacts on the quality. In [2], the authors proposed an algorithm for detecting the freeze frames. The algorithm uses the absolute and squared values of 1-step frame differences, and detects the possible dropped frames by comparing to thresholds that depend on the inherent motion characteristics of the video. In a more recent work [3], the authors also used the squared value of 1-step frame differences, but added an extra encoding pass for the received video. Based on the fact that the frames in a freeze event are exactly the same, the authors' detection algorithm used different thresholds according to the frame types of neighboring frames after that additional encoding. Even though the method of [3] provides more accurate freeze detection than that reported in [1], the required re-encoding process in the method will bring too much overhead to be deployed in a real-time system. In [4], the authors measured the number of freezes as well as the average freeze duration in each preset segment (e.g., the first few seconds, the middle part and the last few seconds) and characterize the quality of each segment by the product of the number of freezes and the average freeze duration. Finally the overall quality of the entire video is modeled using an exponential mapping of a weighted sum of qualities of each segment. The weights depend on the positions of the segments in the video. This metric considers the fact that the degradation caused by frame freeze depends on the freeze locations, but it is still content-independent. In the recent NR metric standardized by ITU-T, namely Rec. ITU-T P.1201.1/.2 [5]–[7], which relies on packet header information, the frame freezing quality degradation is estimated by the ratio of the number of damaged video frames and the total number of video frames as well as the packet loss event frequency. Again, the method did not consider the content characteristic. Furthermore, it does not differentiate between random individual frame drops and consecutive frame losses. Finally this metric involves a series of complicated mapping functions, making it difficult to compute.

With the shortcomings of previous works in mind, our proposed algorithm operates on the decoded video (thus no header information is needed) and it explicitly considers the differences in video content by measuring frame difference statistics. Furthermore, instead of trying to find an analytical mapping between features and the perceived quality, we use a neural network to learn the mapping from the training data. The remainder of the paper is organized as follows: in Section II we will introduce our algorithm for freeze detection and feature extraction. In Section III we report our model selection and training procedure. We show our result in Section IV, and conclude the paper in Section V.

II. Proposed algorithm

Our proposed algorithm consists of three main parts, including detecting freeze frames, extracting features and mapping the features into a quality score. We describe each step in the following sections. The overall flowchart of the proposed metric is illustrated in Fig. 1.

Copyright (c) 2014 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Y. Xue, B. Erkin and Y. Wang are with the Department of Electrical and Computer Engineering, NYU Polytechnic School of Engineering, Brooklyn, NY, 11201 USA. E-mail: {yxue,be521,yw523}@nyu.edu.

TABLE I
DETECTION AND FALSE ALARM OF FREEZE DETECTION ALGORITHM

Dataset	Total freezes	Correctly detected freezes	Detection rate	False alarms	False alarm rate
VQEG	140	131	93.57%	11	7.75%
LIVE	160	160	100.00%	12	6.89%

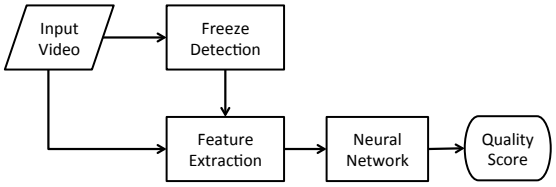

Fig. 1. Flowchart of the proposed metric.

A. Freeze frame detection

Freeze frame detection is a crucial step for our quality metric. A naïve way to finding freeze frames is to examine the difference between each frame and its previous frame, and label the frames that have zero frame difference as freeze frames. This is based on the nature of a freeze event, as the exact same frame is held during the entire freeze duration. In Fig. 2 and 3, we show frame difference profiles from two videos used in this study, where the frame difference ($\mathcal{FD}$) is defined below:

$$\mathcal{FD}(i) = \mathbf{MEAN}((\mathbf{Y}(i+1) - \mathbf{Y}(i))^2) \quad (1)$$

in which $\mathbf{Y}$ is the Y-channel image of a frame. Both Fig. 2 and Fig. 3 have four segments during which the frame difference is zero. They correspond to four freeze events.

Although the zero frame difference appears to be reasonable for detecting frame freezes, in practice, as the quality metric is often computed from the captured video frames rendered on the display screen, there could be small differences between displayed duplicated frames. Therefore, a more robust way to detecting frame freeze is by comparing the frame difference to a non-zero threshold. The challenge lies in how to set the threshold. For example, in the case of videos with large portion of static scenes, even a very small threshold may yield some false alarms. In this work, we use the freeze detection algorithm proposed in [2]. We use the same set of thresholds suggested in [2] for both packet loss and packet delay videos.

In Table I, we report the true freeze detection rate, and false alarm rate for both VQEG and LIVE videos we use in this work. The freeze detection algorithm performs accurately for LIVE videos, while it has some difficulties in picking up freezes for two VQEG videos containing mostly low and smooth motions. We would like to note that the same difficulty would be also applicable to human subjects when the video consists of low and smooth motion. And because of the ambiguity of the freeze and smooth motion in such cases, it is therefore expected that the missed freeze events will not lead to perceivable quality degradation in perceptual quality. We would also like to note that the algorithm will discard any detected freeze event containing less than 2 frames, which is usually not perceivable by the human observer.

B. Feature extraction

After we get the freeze frame locations for a giving video, we extract a total number of 13 features for a video.

First of all, we identify each consecutive set of freeze frames as a freeze event and calculate the duration (in terms of number of frames)

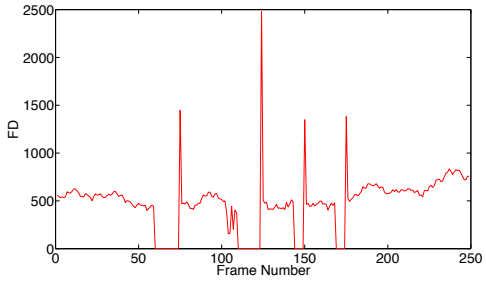

Fig. 2. Frame difference profile for a video suffering from frame freeze due to frame loss, Y channel.

of each freeze event, and the distance (also in terms of number of frames) between every two adjacent freeze events. The number of freezes (**NumFz**) is the number of freeze events for the whole video. **AvgFzDur**, **MaxFzDur**, and **StdFzDur** are the mean, maximum, and standard deviation of the freeze durations among all freeze events. **AvgFzDist**, **MaxFzDist**, and **StdFzDist** are the mean, maximum, and standard deviation of the distances in between two freezes among all freeze events. Finally, **rLenFz** is the ratio of the total freeze length versus the length of the video, and **rDurDist** is the ratio of the average freeze duration and the average inter-freeze distance.

The aforementioned features only depend on the freeze event pattern, but not the actual video content. We also explicitly construct features that are related to the video content, more specifically the frame difference information, which depends on the motion characteristics. The first two features are used to describe a phenomenon typically observed in frame freeze due to packet loss. As shown in Fig. 2, we can see that there are $\mathcal{FD}$ peaks immediately after a freeze due to packet loss. In this paper, we refer this as the "freeze frame difference", denoted by **FzFD**. Intuitively, the frame difference after a freeze is proportional to the length of the freeze, as well as the actual motion during the freeze. Therefore, we define two content-dependent features, denoted by **AvgFzFD** and **MaxFzFD**, which are the mean and maximum of such **FzFD** among all the freeze events for a video.

To compute the background frame difference, we take the average of frame differences over all frames, excluding both the freeze region and scene cut region. In this work, scene cut detection is based on the following heuristic rule: we consider a frame is during a scene cut if the frame difference of the current frame is larger than 5 times the average frame difference of previous five frames.

We name this feature as the background frame difference (**AvgBgFD**). Lastly, the ratio of **AvgFzFD** and **AvgBgFD** is used as another feature (**rFD**). Note that for frame freeze due to packet delay, the frame difference immediately after a freeze will be similar to background frame difference, as illustrated in Fig. 3. Therefore, these features are expected to be less useful for characterizing impact of frame freezes due to delay. In Table II we give a summary of 13 features considered in developing the proposed metric. We would like to note that in a streaming scenario like videoconference, the features can be estimated using the received video. Naïvely, for features involving counting, one could increment the count whenever an event (e.g. freeze) happens; for features related to content, one could use all the received frames up to the last known freeze location to derive the features as the input of the proposed metric. And the received frames can be used to update those estimates periodically or after a freeze completes. The engineering detail of how to implement the proposed metric to real video streaming system, is nevertheless beyond the scope of this paper.

TABLE II
SUMMARY OF 13 FEATURES

Feature name	Description
NumFz	Number of freeze events
AvgFzDur, MaxFzDur, StdFzDur	Average, maximum and standard deviation of freeze durations
AvgFzDist, MaxFzDist, StdFzDist	Average, maximum and standard deviation of distances in between two freezes
rLenFz	Ratio of total freeze length vs. length of the video
rDurDist	Ratio of average freeze duration and the average inter-freeze distance
AvgFzFD, MaxFzFD	Average and maximum post freeze frame difference
AvgBgFD	Average frame difference of video content excluding freezes and scene cuts
rFD	Ratio of **AvgFzFD** and **AvgBgFD**

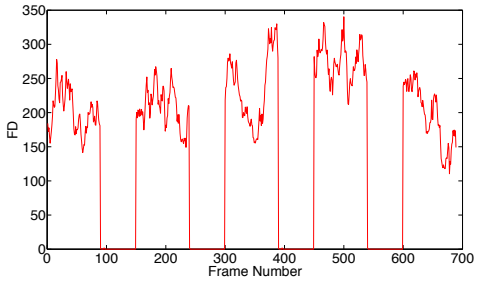

Fig. 3. Frame difference profile for a video suffering from frame freeze due to packet delay, Y channel.

TABLE III
PVS CONFIGURATION FOR VQEGHD5 DATABASE

Video contents	7 source videos, 28 PVSs
Spatial resolution	1080p
Temporal resolution	25Hz
Duration	10 seconds
Encoder	H.264
Bitrate	16Mpbs
PLR	Bursty 0.125, 0.25, 0.5, 1 and 2
Freeze type	Freeze due to packet loss

III. MODEL SELECTION AND TRAINING

A. Dataset

We use two publicly available annotated video databases as our dataset. The VQEGHD dataset 5 [8] contains a total of 28 processed video sequences (PVSs) with frame freezes due to packet loss. The PVS configuration is listed in Table III.

The LIVE mobile FF database [9] contains a total of 40 PVSs, 30 of them have frame freezes due to the packet delay (which have no observed 1-step frame difference spikes immediately after a freeze event, see Fig. 3 for reference), while 10 of them have frame freezes due to packet loss. The PVS configuration for LIVE database is listed in Table IV. We show the screenshots of source videos in Fig. 4. In Fig. 5, we also show a scatter plot of AvgBgFD vs. rFD of all 68 PVs in our database. The scatter plot shows that the considered PVSs cover a large variety of different motion characteristics.

TABLE IV
PVS CONFIGURATION FOR LIVE DATABASE

Video contents	10 source videos, 40 PVSs
Spatial resolution	720p
Temporal resolution	30Hz
Duration	15 seconds
Encoder	RAW
PLR	Every 1, 2 and 4 seconds
Freeze Type	30 PVSs due to packet delay, 10 PVSs due to packet loss.

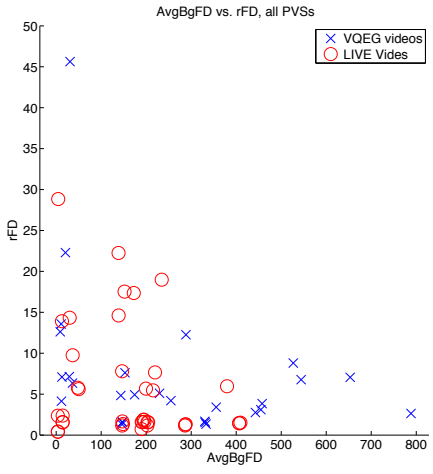

Fig. 5. Scatter plot of AvgBgFD vs. rFD.

We split the whole dataset into a training set and a testing set with $80/20$ proportion, resulting a training set with 52 PVSs from 13 source videos and a testing set with 16 PVSs from 4 source videos. The training set is used to optimize the feature selection and neural network structure and weights. The chosen feature set and network structure and the trained neural network weights are then applied to the testing set to evaluate the performance of the proposed metric. Our testing set only includes source videos not used in the training set. In the testing set, we have a total of 4 source videos (16 PVSs) covering different levels of texture and motion activities. The corresponding snapshots of videos in the testing set are shown in Fig. 4(d)(e)(m)(p). The source videos in both the training set and testing set, respectively, cover typical motion and texture characteristics.

We use the difference mean opinion score (DMOS) as the subjective quality measure for our dataset. The DMOS score is defined as $MOS(PVS) - MOS(SRC) + 5$, where SRC is the hidden reference source.

B. Neural Network Structure Selection

In this work, instead of using a preset function to find the mapping from the extracted features to the quality score, we use neural network based approach. Considering the limited number of training samples, we only consider a neural network with 1 hidden layer with Sigmoid transfer function. Besides the hidden layer neurons, all other neurons (both output neuron and input neurons) are set to be linear. We ran through different number of hidden layer neurons and found that three nodes in the hidden layer gives the best performance (in a scheme described in the following section). Figure 6 shows the network diagram.

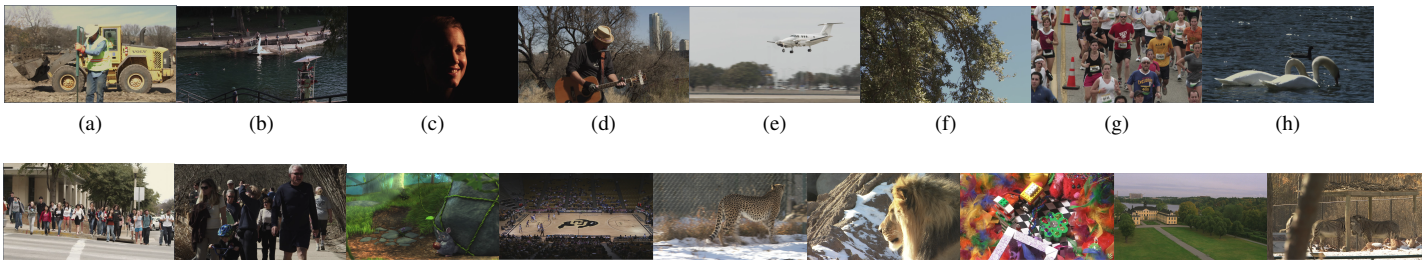

(a) (b) (c) (d) (e) (f) (g) (h)

(i) (j) (k) (l) (m) (n) (o) (p) (q)

Fig. 4. Snapshots for videos used in this work

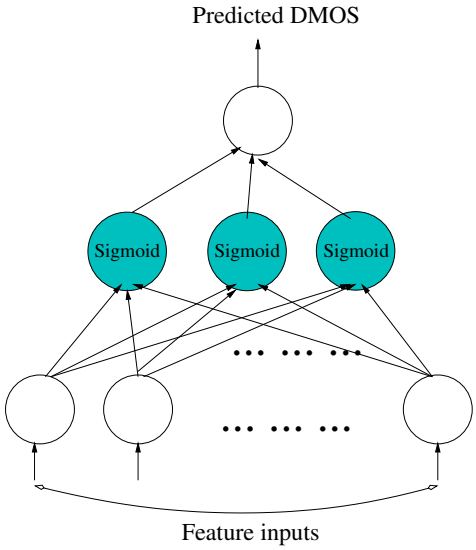

Fig. 6. Neural network architecture, the hidden layer neurons shown in cyan are with sigmoid transfer function.

C. Feature selection and neural network parameter optimization

The one layer neural network can generally accommodate any number of input features, and can have any number of hidden nodes. When the number of input features is $\mathcal{N}$ and the number of hidden nodes is $\mathcal{M}$, the total number of weights to be trained is $\mathcal{M} \times (\mathcal{N}+1) + \mathcal{M} + 1$. Given the limited number of training samples (we have a total of 52 PVSs in training set), we must choose $\mathcal{M}$ and $\mathcal{N}$ so that $\mathcal{M} \times (\mathcal{N}+1) + \mathcal{M} + 1$ is strictly less than 52, to avoid overfitting, and that the $\mathcal{N}$ features should be chosen from all 13 features described in Sec. II-B to maximize a proper performance criterion. We combine our feature selection with the model selection (i.e., $\mathcal{M}$) using the popular 10-fold cross validation procedure [10], that splits the training set into 10 equally numbered batches, each time one batch is used as the validation fold while the rest are used to train the specific network structure of interest. The training procedure is repeated 10 times and the performance is reported as the average error on the validation set.

We use the exhaustive search for selecting both the number of hidden nodes $\mathcal{M}$ and the number of features $\mathcal{N}$. Furthermore, for a given $\mathcal{N}$, we search through all possible feature combinations. We find the combination of $\mathcal{M}$ and $\mathcal{N}$ and corresponding $\mathcal{N}$ features that leads to the minimal cross validation error. Through this procedure we find that $\mathcal{M} = 3$ hidden nodes and $\mathcal{N} = 6$ features give the best performance. Table V lists the chosen six features.

TABLE V
CHOSEN FEATURES THROUGH EXHAUSTIVE SEARCH

Features	AvgFzDist, NumFz, rDurDist, rFD, StdFzDist, rLenFz

TABLE VI
LEARNED WEIGHTS FOR NEURAL NETWORK, WEIGHTS ON HIDDEN NODES ARE ORDERED BASED ON THE FEATURES LISTED IN TABLE V, THE LAST NUMBER FOR EACH NODE IS THE BIAS TERM FOR THAT NODE

Node 1	-0.5236, 2.8352, -0.6619, 2.2123, -0.2637, -0.3205, and 3.1631
Node 2	5.6230, -4.7354, 2.1113, -2.6986, -1.9342, 6.0606, and 1.8050
Node 3	1.6702, -0.8454, 1.8230, -2.4986, 1.5318, -0.2756, and -2.2932
Output layer	-1.2341, -0.5106, -1.1324, and 0.0932

TABLE VII
PERFORMANCE COMPARISON

Metrics	PCC	SROCC	rRMSE
[1]	0.71	0.73	8.87%
[4]	0.54	0.51	12.1%
VQM-VFD [11] on VQEG videos	0.65	0.62	10.1%
Proposed (whole set)	0.93	0.91	5.04%
Proposed (training set)	0.95	0.94	3.34%
Proposed (testing set)	0.80	0.79	6.7%

D. Training and testing

For given $\mathcal{N}$ input features and $\mathcal{M}$ hidden nodes, and for each of the 10 possible training subsets (in the 10-fold cross validation process), we follow the conventional method for training the network weights, to minimize the mean square error between the predicted quality scores and the subjective scores for all training samples. The Jacobian matrix based Levenberg-Marquardt method is used for faster convergence than traditional gradient descent [10]. Once the optimal $\mathcal{M}$ and $\mathcal{N}$ and corresponding feature set is determined, we use all training data (containing 52 samples) to retrain the network weights. The resulting set of weights is listed in Table VI, and is used to predict the quality of the samples in the testing set. All input features are normalized to zero mean and unit variance.

IV. RESULT

Figure 7 shows the scatter plots of subjective ratings (in terms of DMOS) and the predicted scores for the proposed neural network using the optimized network structure and features. Note that the proposed metric achieves a high correlation of over 0.9 for the whole set and 0.8 for a separate testing set. We also show results by the

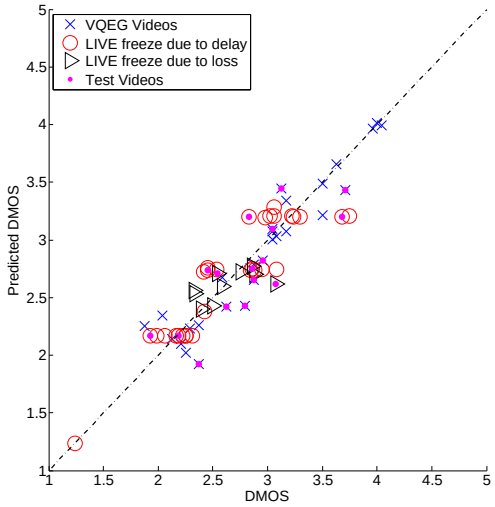

Fig. 7. Scatter plot for the proposed metric using the optimally selected features and the trained neural network structure. Points with "·" are test videos, while others are training videos.

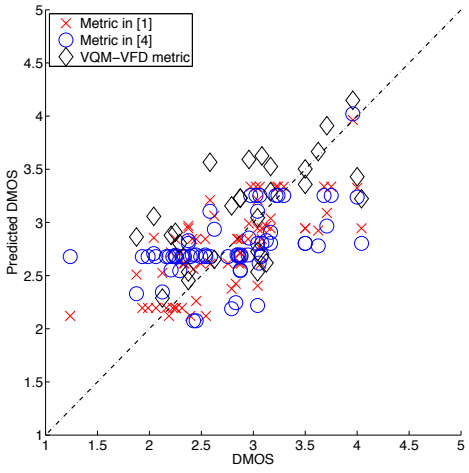

Fig. 8. Scatter plot for NR Metric in [1], NR metric in [4], and Full-referenced VQM−VFD metric in [11]. For [1] and [4] , results for both VQEG and LIVE datasets are reported. For [11], only results are VQEG video are included.

NR temporal jerkiness metric in [1] and [4], and the full reference VQM−VFD [11] metric in Fig. 8.

We report the Pearson correlation coefficient (PCC), Spearman rank order correlation coefficient (SROCC) and relative root-mean-square error (rRMSE) of these methods in Table VII. For [1], we find the optimal parameters using the same training set used in our proposed metric. We have found that the number of bins specified for the histogram is a crucial factor for the metric performance. Because it does not include any video content characteristics, it could not differentiate videos with the same freezing pattern (as is the case for LIVE mobile database). For [4], we used the same definitions and corresponding weights reported for freezes happened in different parts of the video as reported in [4]. For [11], we use the recommended full reference calibration method [12] for temporal alignment. We only report results for the VQEG videos since we are not able to run the provided software [13] on LIVE videos. It is clear that, from both the scatter plots and Table VII, the proposed method correlates more closely to the subjective ratings than the comparison methods. Note that, we are aware that the metric proposed in [4] is designed for longer video sequences, we observe that the limitation of their model is the main reason it performs poorly in comparison. In fact, their metric can be expressed as a one layer neural network.

V. Conclusion

In this paper, we propose a novel NR temporal jerkiness quality metric for videos acting on a carefully chosen set of features. The network structure and the features are optimized through a cross validation procedure using training data extracted from publicly available annotated video datasets.

The resulting feature set and the trained network are further evaluated on separate testing data. The proposed metric works equally well for frame freezes due to either packet loss or packet delay. Our metric extracts 6 features from the distorted video which describe the freeze duration, inter-freeze distance, and the ratio between the frame difference after a freeze and the background frame difference. The last feature depends on the motion characteristics of the video. Our data driven approach then utilizes a 1-hidden layer neural network trained by a subset coming from the publicly available annotated video datasets to find the mapping from the features to the quality. Our trained network with optimally chosen features can achieve very good performance with average PCC over 0.9 for the training set and achieve PCC $= 0.8$ for video contents not seen before. Our proposed metric significantly outperforms two previously reported no-reference video quality metrics and one full reference metric.

References

[1] Q. Huynh-Thu and M. Ghanbari, "No-reference temporal quality metric for video imparied by frame freezing artifacts," in *IEEE International Conference on Image Processing*, 2009, pp. 2221–2224.

[2] S. Wolf, "A no reference (nr) and reduced reference (rr) metric for detecting dropped video frames," in *Int. Workshop on Video Processing and Quality Metrics for Consumer Electronics (VPQM)*, 2009, Scottsdale, Arizona.

[3] G. Yammine, E. Wige, F. Simmet, D. Niederkorn, and A. Kaup, "Blind frame freeze detection in coded videos," in *Picture Coding Symposium*, 2012, Krákow, Poland, pp. 341–344.

[4] D. Z. Rodríguez, J. Abrahão, D. C. Begazo, R. L. Rosa, and G. Bressan, "Quality metric to assess video streaming service over tcp considering temporal location of pauses," *IEEE Transactions on Consumer Electronics*, vol. 58, no. 3, pp. 985–992, 2012.

[5] ITU-T, *ITU-T P.1201.1 Parametric non-intrusive assessment of audiovisual media streaming quality - lower resolution application area*, October 2012.

[6] ——, *Rec. ITU-T P.1201.2 Parametric non-intrusive assessment of audiovisual media streaming quality - higher resolution application area*, October 2012.

[7] K. Yamagishi and G. Shan, "Packet-header-based-quality-estimation model for mobile audiovisual media streaming," in *Int. Workshop on Video Processing and Quality Metrics for Consumer Electronics (VPQM)*, 2013.

[8] A. Webster and F. Speranza, "Report on the validation of video quality models for high definition video content," Video Quality Experts Group, Tech. Rep., July 2010.

[9] A. K. Moorthy, L. K. Choi, A. C. Bovik, and G. de Veciana, "Video quality assessment on mobile devices: subjective, behavioral and objective studies," *IEEE Journal of Selected Topics in Signal Processing*, vol. 6, pp. 652–671, 2012.

[10] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, 2nd ed. Springer, 2009.

[11] S. Wolf and M. H. Pinson, "Video quality model for variable frame delay," NTIA Technical Memorandum, Tech. Rep., September 2011.
[12] M. H. Pinson and S. Wolf, "A new standardized method for objectively measuring video quality," *IEEE Transactions on Broadcasting*, vol. 50, no. 3, pp. 312–322, September 2004.
[13] NTIA. Video quality metric (vqm) software. [Online]. Available: http://www.its.bldrdoc.gov/resources/video-quality-research/request-software.aspx

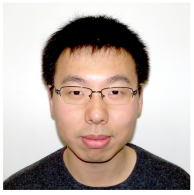

Yuanyi Xue (S09) received the B.Eng.(08) degree in Electrical Engineering from Southeast University, Nanjing, China. He received the M.S. degree in Electrical Engineering from Polytechnic School of Engineering of New York University (then Polytechnic Institute of NYU) in 2010. He is now a Ph.D. candidate major in Electrical Engineering in Polytechnic School of Engineering of New York University. His research interests include perceptual video quality assessment and modeling, sparse signal recovery, and machine learning applications for video signal processing.

Beril Erkin is a Video Software engineer at Dialogic Inc, New Jersey. She received her Bachelors degree with honors in Electrical and Electronics Engineering at Middle East Technical University in Ankara, Turkey, before completing a Masters degree in Electrical Engineering at Polytechnic School of Engineering of New York University in 2014. Her research interests include video communications, WebRTC, image analysis, computer vision, and machine learning.

Yao Wang (M90-SM98-F04) received the B.S. and M.S. degrees in Electronic Engineering from Tsinghua University, Beijing, China, in 1983 and 1985, respectively, and the Ph.D. degree in Electrical and Computer Engineering from University of California at Santa Barbara in 1990. Since 1990, she has been on the faculty of Electrical and Computer Engineering, Polytechnic School of Engineering of New York University (formerly Polytechnic University, Brooklyn, NY). She was on sabbatical leave at Princeton University in 1998 and at Thomson Corporate Research, Princeton, in 2004-2005. She was a consultant with AT& T Labs - Research, formerly AT& T Bell Laboratories, from 1992 to 2000. Her research areas include video communications, multimedia signal processing, and medical imaging. She is the leading author of a textbook titled Video Processing and Communications, and has published over 150 papers in journals and conference proceedings. She has served as an Associate Editor for IEEE TRANSACTIONS ON MULTIMEDIA and IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY. She received New York City Mayors Award for Excellence in Science and Technology in the Young Investigator Category in year 2000. She was elected Fellow of the IEEE in 2004 for contributions to video processing and communications. She is a co-winner of the IEEE Communications Society Leonard G. Abraham Prize Paper Award in the Field of Communications Systems in 2004, and a co-winner of the IEEE Communications Society Multimedia Communication Technical Committee Best Paper Award in 2011. She was a keynote speaker at the 2010 International Packet Video Workshop. She received the Overseas Outstanding Young Investigator Award from the Natural Science Foundation of China in 2005 and was named Yangtze River Lecture Scholar in Tsinghua University by the Ministry of Education of China in 2007.